Preprint ANL/MCS-P3037-0912

On parallelizing dual decomposition in stochastic integer programming

Miles Lubin[1,a], Kipp Martin[b], Cosmin Petra[a], Burhaneddin Sandıkçı[b]

[a] *Mathematics and Computer Science Division, Argonne National Laboratory, Argonne, IL, USA*
[b] *Booth School of Business, University of Chicago, Chicago, IL, USA*

Abstract

For stochastic mixed-integer programs, we revisit the dual decomposition algorithm of Carøe and Schultz from a computational perspective with the aim of its parallelization. We address an important bottleneck of parallel execution by identifying a formulation that permits the parallel solution of the *master* program by using structure-exploiting interior-point solvers. Our results demonstrate the potential for parallel speedup and the importance of regularization (stabilization) in the dual optimization. Load imbalance is identified as a remaining barrier to parallel scalability.

Keywords: stochastic programming, mixed-integer programming, column generation, dual decomposition, parallel computing, bundle methods

1. Introduction

Stochastic mixed-integer programmming (SMIP) models with recourse [1, 2] are commonly used in practice for making discrete decisions under uncertainty. Such models arise in applications in energy, routing, scheduling, production planning, and others, where parts

Email addresses: `mlubin@mit.edu` (Miles Lubin[1]), `kmartin@chicagobooth.edu` (Kipp Martin), `petra@mcs.anl.gov` (Cosmin Petra), `burhan@chicagobooth.edu` (Burhaneddin Sandıkçı)

[1]Present affiliation: MIT Operations Research Center, Cambridge, MA, USA

or all of the data for the model are not completely known at the time decisions must be made, but can be approximated by some known stochastic model.

Although many practical instances remain difficult to solve, significant progress has been made in developing algorithms to solve these problems, particularly those with special structure such as pure-integer recourse or pure-binary first-stage decisions (for reviews, see [1–3]). For more general SMIP problems, Sen [3] suggests the *dual decomposition* (DD) approach of Carøe and Schultz [4] or the *branch-and-price* (BP) approach of Lulli and Sen [5]. This paper focuses on these two approaches from the perspective of parallel computing.

In our theoretical development in §2, we demonstrate an effective equivalence between the nonsmooth Lagrangian dual problem solved by DD and the restricted master problem solved by BP. While it was previously known that these problems have the same optimal values, the effective equivalence is stronger in that solving one provides an optimal solution to both. This fact relates to the so-called primal-recovery properties of subgradient approaches applied to Lagrangian duals, which only recently have become more widely known in the optimization community [6–8]. Both approaches are therefore seen to solve the same relaxation simply by different algorithms for nonsmooth optimization, the former by the proximal bundle method [9] and the latter by a more slowly convergent cutting-plane method, as discussed in §3.

With this equivalence established, we proceed in §4 to present our analysis of and a new formulation for the quadratic program (QP) master problem solved at each iteration of the proximal bundle method, considering the particular structure induced by relaxing nonanticipativity constraints. The new formulation results in a block-angular QP that can be solved efficiently by recently developed interior-point solvers for structured QPs. Improvements of several orders of magnitude are observed over the QP solvers from off-the-shelf proximal bundle codes.

Reducing the time spent solving the master problem significantly increases the scope for parallelism, which has previously been identified but not exploited in an implementation. In §5, we present our numerical results from a preliminary parallel implementation on a high-performance cluster.

2. Dual decomposition and branch-and-price

Consider the following two-stage SMIP with recourse:

$$z = \min\{c^\top x + Q(x) : \quad Ax \le b, x \in X\}, \tag{1}$$

where

$$Q(x) = \mathbb{E}_\xi \left[\min\left\{q(\xi)^\top y : \quad Wy \le h(\xi) - T(\xi)x, y \in Y\right\}\right]. \tag{2}$$

The parameters $c \in \mathbb{R}^{n_1}$, $b \in \mathbb{R}^{m_1}$, $A \in \mathbb{R}^{m_1 \times n_1}$, and $W \in \mathbb{R}^{m_2 \times n_2}$ are fixed and known. The vector ξ is a random variable, which we assume here to have a discrete distribution with r possible realizations $\xi_1, \ldots, \xi_r$ and corresponding probabilities $p_1, \ldots, p_r$. Realization $j = 1, \ldots, r$, known as *scenario* j, contains the data $(q(\xi_j), h(\xi_j), T(\xi_j))$, now (q_j, h_j, T_j) for brevity, where the vectors q_j and h_j and the matrix T_j have conformable dimensions. The sets $X \subseteq \mathbb{R}^{n_1}_+$ and $Y \subseteq \mathbb{R}^{n_2}_+$ denote restrictions that some or all of the variables take integer or binary values. For $j = 1, \ldots, r$, define the set

$$S_j := \{(x, y_j) : \quad Ax \le b, x \in X, T_j x + W y_j \le h_j, y_j \in Y\}.$$

The *deterministic equivalent* problem to (1), which we assume to be feasible and bounded, is

$$z = \min\left\{c^\top x + \sum_{j=1}^{r} p_j q_j^\top y_j : \quad (x, y_j) \in S_j, j = 1, \ldots, r\right\}. \tag{3}$$

Also consider the equivalent *split-variable* formulation [10]:

$$z = \min\left\{ \sum_{j=1}^{r} p_j(c^\top x_j + q_j^\top y_j) : \quad (x_j, y_j) \in S_j, \quad j = 1, \ldots, r, \quad x_{\cdot} = x_1 = \ldots = x_r \right\}. \quad (4)$$

The constraints $x_{\cdot} = x_1 = \ldots = x_r$ are known as the nonanticipativity conditions, which force the first-stage decision x to be the same under each scenario. We have introduced an additional variable $x_{\cdot}$ and the nonanticipativity constraints are represented as

$$x_j - x_{\cdot} = 0, \quad j = 1, \ldots, r. \quad (5)$$

This representation of non-anticipativity, as used by Lulli and Sen [5] differs from the one used by Carøe and Schultz [4], who instead represent it by a set of equalities solely on the variables $x_1, \ldots, x_r$ of the form $\sum_{j=1}^{r} H_j x_j = 0$. (For example, $x_1 - x_j = 0, \ j = 2, \ldots, r$.) The representations are equivalent; however, we see later that the form used here is advantageous for computation.

Relaxing the nonanticipativity constraints (5), one may write the Lagrangian relaxation of (4) as

$$D(\lambda_1, \ldots, \lambda_r) = \min\left\{ \sum_{j=1}^{r} \left[L_j(x_j, y_j, \lambda_j) - \lambda_j^\top x_{\cdot} \right] : \quad (x_j, y_j) \in S_j, \quad j = 1, \ldots, r \right\}, \quad (6)$$

where $L_j(x_j, y_j, \lambda_j) = p_j(c^\top x_j + q_j^\top y_j) + \lambda_j^\top x_j$ for $j = 1, \ldots, r$. As $x_{\cdot}$ is unconstrained, the condition $\sum_{j=1}^{r} \lambda_j = 0$ is required for boundedness of (6). With this condition, the $\lambda_j^\top x_{\cdot}$ terms vanish, and (6) is separable into $D(\lambda_1, \ldots, \lambda_r) = \sum_{j=1}^{r} D_j(\lambda_j)$, where, for $j = 1, \ldots, r$,

$$D_j(\lambda_j) = \min_{x_j, y_j} \left\{ L_j(x_j, y_j, \lambda_j) : \quad (x_j, y_j) \in S_j \right\}. \quad (7)$$

For any choice of $\lambda_1, \ldots, \lambda_r$, it is clear that $D(\lambda_1, \ldots, \lambda_r) \le z$; that is, the Lagrangian relaxation provides a valid lower bound on the optimal value z of (1). A natural problem is then to find the best such bound. This is known as the Lagrangian dual problem, which is expressed as [3]

$$z_{\mathrm{LD}} = \max_{\lambda_1,\ldots,\lambda_r} \left\{ \sum_{j=1}^{r} D_j(\lambda_j) : \sum_{j=1}^{r} \lambda_j = 0 \right\}. \tag{8}$$

Because of the nonconvexity introduced by the integer requirements, the optimal value z_{LD} is typically, but not always, strictly less than z. We restate Proposition 2 of [4], which provides a characterization of the optimal value.

Proposition 1. *The optimal value z_{LD} of the Lagrangian dual* (8) *equals the optimal value of the linear program*

$$\min \left\{ \sum_{j=1}^{r} p_j (c^\top x_j + q_j^\top y_j) : \quad (x_j, y_j) \in \operatorname{conv}(S_j)\,, \quad x_j = x_\cdot, \quad j = 1, \ldots, r \right\}. \tag{9}$$

The Lagrangian dual (8) is a concave, nonsmooth optimization problem, which Carøe and Schultz [4] propose to solve with subgradient methods (or more properly in this context, supergradient methods). The bounds generated from the Lagrangian dual are used within a branch-and-bound procedure. This is the so-called dual decomposition (DD) approach.

On the other hand, Lulli and Sen [5] propose to solve (9) directly using a *column generation* procedure. Below we show that these two approaches are duals of each other, and that it is easy to recover a primal solution to (9) from DD. This primal solution can then be used as in a branch and price algorithm.

In the rest of this section, we first present the classical subgradient cutting-plane approach and then demonstrate that applying the cutting-plane method to (8) is equivalent to solving (9) by column generation.

Observe that (8) is equivalent to

$$\begin{aligned} \max \quad & \sum_{j=1}^{r} \theta_j & & \qquad (10) \\ \text{s.t.} \quad & \sum_{j=1}^{r} \lambda_j & = \;& 0, \qquad (11) \\ & \theta_j & \leq \;& D_j(\lambda_j), \quad j = 1, \ldots, r. \qquad (12) \end{aligned}$$

Each $D_j(\lambda_j)$ is concave in λ_j, and we say that γ_j^k is a subgradient of $D_j(\lambda_j)$ at the point λ_j^k if, for all λ_j,

$$D_j(\lambda_j) \leq D_j(\lambda_j^k) + (\gamma_j^k)^\top (\lambda_j - \lambda_j^k).$$

Since $L_j(x_j, y_j, \lambda_j) = p_j(c^\top x_j + q_j^\top y_j) + \lambda_j^\top x_j$, given a λ_j^k, the corresponding subgradient γ_j^k is equal to x_j^k, where (x_j^k, y_j^k) is a solution to (7). The cutting-plane method (as depicted by the pseudocode in Figure 1) replaces each $D_j(\lambda_j)$ in (12) with a relaxation using a set of subgradients and solves the following linear program at each iteration:

$$\begin{aligned} \max \quad & \sum_{j=1}^{r} \theta_j & & \qquad (13) \\ \text{s.t.} \quad & \sum_{j=1}^{r} \lambda_j & = \;& 0, \qquad (14) \\ & \theta_j & \leq \;& D_j(\lambda_j^k) + (x_j^k)^\top (\lambda_j - \lambda_j^k), \quad j = 1, \ldots, r, \quad k = 1, \ldots, K. \qquad (15) \end{aligned}$$

```
Initialize:  Choose a relative convergence tolerance $\epsilon$.
             $K \leftarrow 1$, $\lambda_j^K \leftarrow 0$ for $j = 1, \ldots, r$.
             Solve (7) for $j = 1, \ldots, r$, saving optimal value $D_j(\lambda_j^K)$ and solution $x_j^K$.
Step 1:      Solve (13)-(15), obtaining optimal $\theta_j^*$ and $\lambda_j^*$ for $j = 1, \ldots, r$.
Step 2:      $K \leftarrow K + 1$, $\lambda_j^K \leftarrow \lambda_j^*$ for $j = 1, \ldots, r$.
             Solve (7) for $j = 1, \ldots, r$, saving optimal value $D_j(\lambda_j^K)$ and solution $x_j^K$.
Step 3:      If $\sum_j \big[\theta_j^* - D_j(\lambda_j^K)\big] / \big[1 + |\sum_j D_j(\lambda_j^K)|\big] < \epsilon$ terminate;
             else add $D_j(\lambda_j^K) + (x_j^K)^\top (\lambda_j - \lambda_j^K)$ to (15).
Step 4:      Goto Step 1
```

Figure 1: Pseudocode for cutting-plane algorithm.

It is easy to recover a primal solution to (9) from the linear programming solution of (13)-(15). Assign dual variables $x_{\cdot}$ to (14) and z_j^k to (15). The dual of (13)-(15) is

$$\min \quad \sum_{j=1}^{r}\sum_{k=1}^{K}\left[D_j(\lambda_j^k) - (x_j^k)^\top \lambda_j^k\right] z_j^k \tag{16}$$

$$\text{s.t.} \quad \sum_{k=1}^{K} z_j^k = 1, \quad j = 1, \ldots, r, \tag{17}$$

$$\sum_{k=1}^{K} z_j^k x_j^k = x_{\cdot}, \quad j = 1, \ldots, r, \tag{18}$$

$$z_j^k \geq 0, \quad j = 1, \ldots, r, \quad k = 1, \ldots, K. \tag{19}$$

For a given λ_j^k, let (x_j^k, y_j^k) be the corresponding optimizer in (7). Minimizing a linear function over the feasible region S_j, of a mixed-integer linear program, is equivalent to minimizing over the convex hull $\text{conv}(S_j)$, of the feasible region. Therefore, $D_j(\lambda_j) = \min_{x_j,y_j}\{L_j(x_j, y_j, \lambda_j) : \ (x_j, y_j) \in \text{conv}\,(S_j)\}$. Then from the definition of $D_j(\lambda_j^k)$ and $L_j(x_j^k, y_j^k, \lambda_j^k)$, the objective function (16) for the column generation problem is

$$\begin{aligned}\min \sum_{j=1}^{r}\sum_{k=1}^{K}\left[D_j(\lambda_j^k) - (x_j^k)^\top \lambda_j^k\right] z_j^k &= \sum_{j=1}^{r}\sum_{k=1}^{K}\left[(p_j(c^\top x_j^k + q_j^\top y_j^k) + (x_j^k)^\top \lambda_j^k) - (x_j^k)^\top \lambda_j^k\right] z_j^k, \\ &= \sum_{j=1}^{r}\sum_{k=1}^{K} p_j(c^\top x_j^k + q_j^\top y_j^k) z_j^k,\end{aligned}$$

and we have the standard restricted master for (9), as described in Lulli and Sen [5]. Hence, feasible (and at convergence, optimal) solutions to (9) are obtained from the dual solution of (13)-(15).

Indeed, by using nearly any subgradient-based method to solve (8), one may obtain at a minimal cost an optimal solution to (9). This fact, while well known to specialists in the general case (see, e.g., [7, 11, 12]), has only recently come to attention in the context of Lagrangian relaxation in integer programming [6, 8]. To the best of our knowledge, this result has not been stated in the context of DD and BP.

The branch-and-bound algorithm used in DD calls for branching on *disagreements* in the primal solutions x_j^k, $j = 1, \ldots, r$, produced by the subproblems, whereas in BP the solution to (9) is used for branching decisions, similar to how the solution to the LP relaxation is used in classical branch and bound for integer programs. Hence, while known to use the same relaxation (9) in a theoretical sense, DD and BP have been viewed as computationally different approaches [3, 13]; we have demonstrated a closer computational connection than previously observed. Based on the effective equivalence between Lagrangian relaxation and column generation, Frangioni [6] suggests the potential for using both the solution to the convexification (9) and the primal solutions to the subproblems within branch and bound. We leave the exploration of this possibility in the present context for future research.

3. Improvements to the cutting-plane algorithm

The cutting-plane algorithm, which in the previous section was shown to be computationally equivalent to column generation, is known to be unstable and to converge slowly on practical instances [7, 9]. Modern algorithms for nonsmooth optimization typically apply some form of regularization to the standard cutting-plane approach, potentially resulting in a more difficult master program but also providing a significant reduction in the total number of iterations required. In particular, the proximal bundle method [9] uses a quadratic penalty in the objective to indirectly regulate the step length at each iteration. This approach has appeared in the stochastic programming literature as Ruszczyński's regularized decomposition [14]. Other approaches include the ℓ_∞ trust-region approach, also known as the boxstep method [15], and level regularization [16, 17]; these approaches have been used in the context of stochastic programming, for example, by [18] and [19, 20], respectively. The relative performance of different forms of regularization is generally not well understood and is typically problem dependent [21, 22].

We focus on the proximal bundle method, which is the most widely used regularization method. In this variant, a quadratic penalty term $\sum_{j=1}^{r} ||\lambda_j - \lambda_j^+||_2^2$ is subtracted from the objective function (13), and the modified master (13)-(15) is

$$\begin{aligned}
\max_{\theta,\lambda} \quad & \sum_{j=1}^{r} \theta_j - \frac{1}{2}\tau \sum_{j=1}^{r} ||\lambda_j - \lambda_j^+||_2^2 && (20)\\
\text{s.t.} \quad & \sum_{j=1}^{r} \lambda_j = 0, \\
& \theta_j \leq D_j(\lambda_j^k) + (x_j^k)^\top(\lambda_j - \lambda_j^k), \quad j = 1, \dots, r, \quad k = 1, \dots, K,
\end{aligned}$$

where $(\lambda_1^+, \lambda_2^+, \dots, \lambda_r^+)$ is the current "prox-center" with $\sum_{j=1}^{r} \lambda_j^+ = 0$. The regularization parameter τ is typically adjusted at each iteration; see [23]. Let $\beta_j := \lambda_j - \lambda_j^+$, and consider the reformulation

$$\begin{aligned}
\max_{\theta,\beta} \quad & \sum_{j=1}^{r} \theta_j - \frac{1}{2}\tau \sum_{j=1}^{r} ||\beta_j||_2^2 && (21)\\
\text{s.t.} \quad & \sum_{j=1}^{r} \beta_j = 0 && (w)\\
& \theta_j - (x_j^k)^\top \beta_j \leq D_j(\lambda_j^k) + (x_j^k)^\top(\lambda_j^+ - \lambda_j^k), \quad j = 1, \dots, r, \quad k = 1, \dots, K. && (z_j^k)
\end{aligned}$$

It is typically advantageous to solve the Lagrangian dual of (21):

$$\begin{aligned}
\min_{w,z} \quad & \sum_{j=1}^{r} \left(\sum_{k=1}^{K} z_j^k \left(D_j(\lambda_j^k) + (x_j^k)^T(\lambda_j^+ - \lambda_j^k) \right) + \frac{1}{2\tau} ||w - \sum_{k=1}^{K} z_j^k x_j^k||^2 \right) && (22)\\
\text{s.t.} \quad & \sum_{k=1}^{K} z_j^k = 1, \quad j = 1, \dots, r, \\
& z_j^k \geq 0, \qquad j = 1, \dots, r, \quad k = 1, \dots, K.
\end{aligned}$$

Using the solution of (22), one can recover the optimal β_j by $\beta_j = \frac{1}{\tau}\left(\sum_{k=1}^{K} z_j^k x_j^k - w\right)$. Since w is an n_1 component vector and the z_j^k are scalars, (22) has $K \times r + n_1$ variables, which is typically a significantly smaller number than the $(n_1+1)\times r$ variables of (13)-(15) or (20). We also expect $K \times r >> n_1$. If this does not hold, it could be advantageous to eliminate w by noting that $w = \frac{1}{r}\sum_{j=1}^{r}\sum_{k=1}^{K} z_j^k x_j^k$ at optimality (this may be derived from the Karush-Kuhn-Tucker conditions). However, this elimination destroys the particular structure that is discussed in §4, and a general sparsity-exploiting solver can perform this elimination automatically.

From standard convergence results [9], we have $||w - \sum_{k=1}^{K} z_j^k x_j^k||^2 \to 0$, $j = 1,\ldots,r$ at convergence of the proximal bundle method. Hence, in the limit, $w = \sum_{k=1}^{K} z_1^k x_1^k = \cdots = \sum_{k=1}^{K} z_r^k x_r^k$, which are precisely the constraints (18) of the standard column generation master. So, one recovers the solution to the convexification (9) directly as the optimal w at convergence of the proximal bundle method.

4. Parallel solution of the master program

It has been observed (e.g., in [5]) that the cutting-plane algorithm described in Figure 1 exhibits scope for parallelism in Step 2, where r independent integer programs must be solved. The same observation holds for regularized variants discussed in §3. The potential for parallel speedup, however, is limited according to Amdahl's law [24] by the serial execution bottleneck of solving the master program, for example, the LP (13)-(15) or the QP (22). In this section, we address this bottleneck by identifying the scope for parallelism in solving the master itself.

A key observation is that the cutting-plane master (14)-(15) has a primal block-angular structure; that is, its constraint matrix can be permuted to the form

$$\begin{pmatrix} X & X & \cdots & X \\ X & & & \\ & X & & \\ & & \ddots & \\ & & & X \end{pmatrix}. \tag{23}$$

The linking rows correspond to the constraints (14), and the diagonal blocks correspond to the constraints (15) for each j. This property is a direct result of the equality-constrained formulation (8); in particular, it does not hold if the nonanticipativity representation of Carøe and Schultz [4] is used.

The proximal bundle master also has this structure in its primal form (21); the computational form (22) has a dual block-angular structure (with constraints in the form of the transpose of (23)). Note that for quadratic programs to be considered to have block-angular structure, the Hessian matrix must additionally be permutable to the following form:

$$\begin{pmatrix} X & X & \cdots & X \\ X & X & & \\ \vdots & & \ddots & \\ X & & & X \end{pmatrix}. \tag{24}$$

Fortunately this structure also holds for (22), because no quadratic terms link z_j^k with $z_{j'}^{k'}$ if $j \neq j'$ for any k, k'.

Block-angular structure in linear and quadratic programs has been successfully exploited for parallelization within interior-point methods [25, 26]. We follow this approach, a discussion of which is beyond the scope of this paper. Only minimal development, if any, is required to efficiently solve (22) using an existing structure-exploiting interior-point code.

We note that Kiwiel [27] developed a specialized active-set method for solving the QP master of the proximal bundle method for unstructured problems. However, this active-set method cannot immediately accommodate the equality constraints of our formulation and it is unknown whether this approach could be successfully parallelized for block-angular structure.

5. Implementation and numerical results

In this section, we explore different computational aspects of dual decomposition, with a view toward parallel computation. All experiments were performed on *Fusion*, a 320-node computing cluster at Argonne National Laboratory. Fusion has an InfiniBand QDR interconnect, and each node has two 2.6 GHz Xeon processors (total 8 cores) and 36 GB of RAM. Serial experiments were performed on a single node of *Fusion*.

We used publicly available two-stage SMIP instances. The `dcap` and `sslp` instances are available at `http://www2.isye.gatech.edu/~sahmed/siplib/`, and the `prod` instances are available at `http://people.orie.cornell.edu/huseyin/research/sp_datasets/sp_datasets.html`. Basic statistics about these instances are listed in Table 1. We refer the reader to the indicated websites for further descriptions of the instances. Because of space limitations, we report only on a subset of the instances available online.

5.1. Serial experiments

We use the `sslp` and `dcap` instances to compare the performance of various methods, executed in serial, for optimizing the Lagrangian dual (8). We experiment with three methods: (*i*) the classical cutting-plane method (Figure 1), (*ii*) the proximal bundle method (Figure 2), and (*iii*) the ℓ_∞ trust-region (boxstep) method using the trust-region updating rules of Linderoth and Wright [18]. All algorithms use a relative convergence tolerance $\epsilon = 10^{-7}$, although the convergence criteria have slight mathematical differences. In our C++ implementation of the cutting-plane and ℓ_∞ trust-region methods, the master program, which is linear, is solved by `Clp` [28], hot-started by using the optimal basis from the

Table 1: Test problem statistics. `prod` instances do not have integer restrictions.

Test	1st Stage			2nd-Stage Scenario		
Problem	Vars.	Intgr.	Cons.	Vars.	Intgr.	Cons.
`sslp_5_25`	5	5	1	130	125	30
`sslp_10_50`	10	10	1	510	500	60
`sslp_15_45`	15	15	1	690	675	60
`dcap233`	12	6	6	27	27	15
`dcap243`	12	6	6	36	36	18
`dcap332`	12	6	6	24	24	12
`dcap342`	12	6	6	32	32	14
`prod-small`	50	0	10	250	0	220
`prod-medium`	250	0	10	1400	0	250
`prod-large`	1,500	0	75	1,450	0	700

previous solution. For the proximal bundle method, we use our implementation as well as the off-the-shelf open-source implementation `ConicBundle` [29].

Within our implementation, we experiment with solving the QP (22) using a general sparsity-exploiting interior-point solver OOQP [30] (with the MA57 [31] sparse linear-algebra routines) and then using the block-angular-structure-exploiting interior-point solver PIPS-IPM [32] (with the LAPACK [33] routines for dense linear algebra). Both OOQP and PIPS-IPM use the same algorithmic implementation of Mehrotra's predictor-corrector scheme [34], and each instance is solved from scratch. All mixed-integer subproblems are solved by using the software package SCIP [35].

Results are presented in Table 2. For each instance and solution method, we report the total number of iterations, the total execution time, the time spent solving the master program, and the objective value at convergence. We first observe that the cutting-plane method requires the largest number of iterations, as expected. For the `sslp` instances, the proximal bundle method is superior (in terms of total execution time) to the ℓ_∞ trust-region approach, while the trust-region approach appears to be superior on the `dcap` instances. We note the disagreement on objective values, in particular for the `dcap` instances. This is explained partially by differing convergence criteria, although in some cases numerical

Table 2: Summary of result with serial experiments. All methods except `ConicBundle` were implemented by the authors. The cutting-plane method is shown to require many more iterations than regularized variants. Asterisk indicates exceeded time limit (7,200 seconds).

Instance			Time (Sec.)		
(Scenarios)	Method	Iter.	Total	Master	Objective
`sslp_5_25` (50)	Cutting plane	46	172	0.09	-121.6
	ℓ_∞ trust region	21	71	0.05	-121.6
	`ConicBundle`	15	55	0.38	-121.6
	OOQP	9	29	0.06	-121.6
	PIPS-IPM	9	29	0.10	-121.6
`sslp_10_50` (50)	Cutting plane	73	2508	0.90	-364.64
	ℓ_∞ trust region	71	4352	0.90	-364.64
	`ConicBundle`	27	1521	4.82	-364.64
	OOQP	14	571	0.18	-365.62
	PIPS-IPM	22	1004	0.44	-364.41
`sslp_15_45` (10)	Cutting plane	89	5088	0.20	-260.5
	ℓ_∞ trust region	65	5762	0.11	-260.5
	`ConicBundle`	36	1628	0.14	-260.5
	OOQP	39	2374	0.48	-260.5
	PIPS-IPM	38	2408	0.26	-260.5
`dcap233` (200)	Cutting plane	194	1189	116	1837.87
	ℓ_∞ trust region	58	295	46	1833.37
	`ConicBundle`	34	*7200	7027	*
	OOQP	58	337	69	1833.4
	PIPS-IPM	68	341	34	1833.4
`dcap243` (200)	Cutting plane	252	1920	168	2326.13
	ℓ_∞ trust region	40	189	25	2322.37
	`ConicBundle`	34	*7200	7048	*
	OOQP	68	348	106	2321.21
	PIPS-IPM	68	266	32	2321.21
`dcap332` (200)	Cutting plane	321	1348	189	1059.10
	ℓ_∞ trust region	47	210	74	1059.08
	`ConicBundle`	33	*7200	7079	*
	OOQP	77	363	121	1059.08
	PIPS-IPM	79	282	43	1059.10

instability is also present; for example, the cutting-plane method for `dcap233` reports a mathematically invalid objective value that is larger than the optimal objective value of the original stochastic integer problem.

For the `sslp` instances, our implementation of the proximal bundle method has a comparable iteration count to that of the `ConicBundle` package, empirically confirming our algorithmic implementation. However, `ConicBundle` is unable to solve the dcap instances to completion because of the time spent solving the QP master. `ConicBundle` uses a similar interior-point method to that of our implementation; the difference in execution time is attributable to its use of dense linear algebra within the interior-point method. That is, `ConicBundle` treats the Hessian matrix of the QP master as entirely dense, whereas in Section 4 it was shown to be highly structured. (Note that `ConicBundle` does not solve the equality-constrainted formulation; the Hessian, however, remains highly structured.) The results demonstrate that the general sparse linear algebra routines used within OOQP (which *could* be implemented within `ConicBundle`) significantly reduce the computation time. In addition, specialized linear algebra for the block-angular structure (as used within PIPS-IPM) can produce a further improvement even before considering parallel computation.

5.2. Parallel experiments

A preliminary parallel version of the proximal bundle method (Figure 2) was implemented by using the Message Passing Interface (MPI) API [36]. In our implementation, each scenario is statically assigned to an *MPI process*, which may be considered a parallel worker. This worker is then responsible for solving the mixed-integer subproblems for its assigned scenarios at each iteration. This static assignment is simpler to implement but is expected to be inferior in its load-balancing properties to a scheme where workers are dynamically assigned to subproblems.

Table 3 contains results from parallel experiments with instances similar to those of the serial experiments in Table 2 but with larger numbers of scenarios. Unlike in Table 2, we

```
Initialize: Choose a relative convergence tolerance $\epsilon$.
            $K \leftarrow 1$, $\lambda_j^+ \leftarrow 0$, $\tau \leftarrow 1$, $m = 0.1$.
            Solve (7) with $\lambda_j^+$ for $j = 1, \ldots, r$, saving optimal solution $x_j^K$.
            $curObj \leftarrow \sum_j D_j(\lambda_j^+)$.
Step 1:     Solve (22), obtaining optimal $w^*, z_j^{k*}, \theta_j^*, \lambda_j^*$.
Step 2:     Let $v = (\sum_i \theta_i^*) - curObj$. If $v/(1 + |curObj|) < \epsilon$, terminate, else continue.
Step 3:     $K \leftarrow K + 1$.
Step 4:     Solve $D_j(\lambda_j^*)$ for $j = 1, \ldots, r$, saving optimal value $D_j(\lambda_j^K)$ and solution $x_j^K$.
Step 5:     $newObj \leftarrow \sum_j D_j(\lambda_j^K)$. Let $u = 2\tau(1 - (newObj - curObj)/v)$.
Step 6:     Update $\tau \leftarrow \min(\max(u, \tau/10, 10^{-4}), 10\tau)$. (See [11])
Step 7:     If $(newObj - curObj > m \cdot v)$ update $\lambda_j^+ \leftarrow \lambda_j^*$, $curObj \leftarrow newObj$.
Step 8:     Goto Step 1.
```

Figure 2: Pseudocode of proximal bundle method as implemented. Note that we take a maximization view, so some signs are flipped from typical statements of the algorithm, e.g., in [11]. An important mathematical feature of bundle methods is the ability to remove old subgradients ("compress the bundle") after Step 4; however, we do not implement this.

consider only the proximal bundle method. Again we solve the master QP using the general sparsity-exploiting QP solver OOQP and then using the structure-exploiting solver PIPS-IPM. The structure-exploiting QP solver is run in parallel using the same MPI processes as the mixed-integer subproblems. Identical runs were performed with 1, 8, 16, and 32 parallel processes, each corresponding to a physical processing core. Recall that the nodes of the compute cluster have 8 cores each; hence 32 processes corresponds to 4 physical nodes.

We note a wide range of behavior on the six instances considered. On all instances except `sslp_5_25`, we observe significant speedups in the time to solve the master QP by using PIPS-IPM on up to 16 parallel processes, although the impact on the total execution time due to the speedups in solving the master varies from the `dcap` instances to the `sslp` instances. For the `dcap` instances, solving the master QP forms a significant portion of the execution time in serial. Hence, by solving the master in parallel, in addition to the mixed-integer subproblems, significant reductions in the total execution time are observed. For `sslp_10_50`, the execution time is dominated by the mixed-integer subproblems, and so speedups in the master have little effect.

Perhaps the most surprising result is the lack of speedup in solving the mixed-integer subproblems when more parallel processors are used. For example, there is little speedup on `sslp_10_50` from 16 to 32 processes. In some cases the total time *increased* from 16 to 32 processes. These results can be explained by the high variability in the time to solve the subproblems as well as by the lack of dynamic load balancing in our implementation. Although the behavior can be explained in retrospect, the solution of the subproblems is typically expected to be a "trivially" parallel computation, yet here it is seen to be far from such. Further work, both computational and theoretical, will be required to address this issue.

Returning to the scalability of the master QP, we conducted an experiment evaluating the relative performances of OOQP and PIPS-IPM on instances with larger numbers of first-stage variables. As noted in §3, formulation (22) may not be efficient unless $K \times r >> n_1$, recalling that n_1 is the number of first-stage variables. Both the `sslp` and `dcap` instances have a small number of first-stage variables. Because we are not aware of SMIP instances with a larger number of first-stage variables, we use the linear `prod` instances with 1,000 scenarios (generated by simple Monte Carlo sampling). This substitution is valid because the structure of the QP master remains the same, and we do not consider the time spent in the subproblems (which are now linear programs). The results in Figure 3 demonstrate that as the number of first-stage variables increases, the general sparse QP solver may become more effective than structure-exploiting solver in serial; yet, when run in parallel, the structure-exploiting solver is significantly faster.

6. Conclusions

This work identifies dual decomposition as a scalable approach for solving stochastic mixed-integer programs, a class of optimization problems known for their computational difficulty. We present the first set of results with a parallel implementation and have addressed the serial bottleneck of solving the master program. Further work, both theoretical

Table 3: Parallel experiments with the proximal bundle method. With OOQP, only the MIP subproblems are solved in parallel; with PIPS-IPM, both the MIP subproblems and the master QP are solved in parallel. For the `dcap` instances, significant overall speedups are observed as a result of reducing the time spent solving the QP master.

		Serial Sparse QP solver (OOQP)				Parallel QP Solver (PIPS-IPM)			
Instance	Parallel		Time (Sec.)				Time (Sec.)		
(Scenarios)	Processes	Iter.	Total	Master	Objective	Iter.	Total	Master	Objective
`sslp_5_25` (100)	1	8	50.5	0.10	-127.370	8	50.6	0.18	-127.370
	8	8	9.1	0.09	-127.370	8	9.7	0.51	-127.370
	16	8	5.7	0.10	-127.370	8	5.9	0.06	-127.370
	32	8	4.3	0.10	-127.370	8	4.2	0.06	-127.370
`sslp_10_50` (500)	1	26	85,837	11.5	-349.132	22	64,659	5.2	-349.133
	8	31	38,826	15.4	-349.132	24	25,178	1.9	-349.136
	16	27	33,060	11.6	-349.131	28	28,142	1.1	-349.136
	32	31	34,274	16.4	-349.137	27	23,349	1.3	-349.118
`dcap233` (500)	1	68	1,098	308	1,736.678	66	839	85.3	1,736.674
	8	68	450	297	1,736.678	70	167	14.0	1,736.681
	16	68	391	298	1,736.678	73	114	9.5	1,736.681
	32	68	674	296	1,736.678	70	87	8.4	1,736.674
`dcap243` (500)	1	57	819	174	2,165.479	57	690	55.6	2,165.479
	8	57	287	169	2,165.479	58	123	8.8	2,165.492
	16	57	228	169	2,165.479	59	122	5.4	2,165.490
	32	57	414	168	2,165.479	59	111	6.0	2,165.495
`dcap332` (500)	1	82	1108	413	1,587.435	80	756	127.2	1,587.256
	8	82	545	407	1,587.435	79	134	15.8	1,587.391
	16	82	476	408	1,587.435	80	151	9.7	1,587.123
	32	82	918	406	1,587.435	77	110	8.1	1,587.439
`dcap342` (500)	1	59	872	163	1,902.842	71	857	89.6	1,903.014
	8	59	356	159	1,902.842	67	214	10.4	1,903.214
	16	59	322	160	1,902.842	56	155	4.3	1,902.893
	32	59	475	159	1,902.842	62	161	5.2	1,902.894

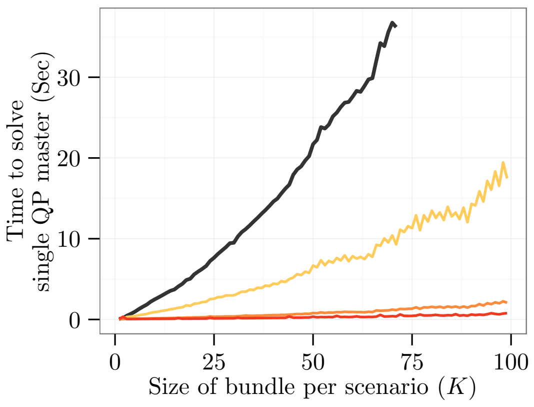

(a) $n_1 = 50$ first-stage variables

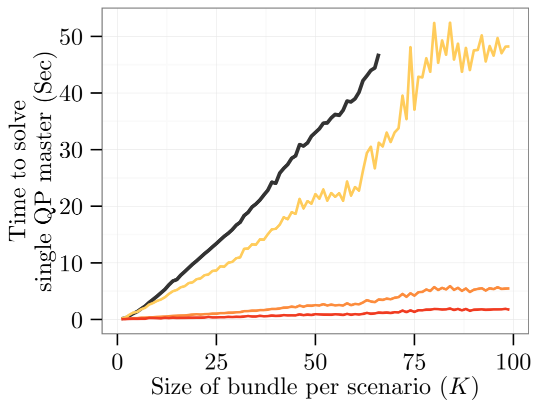

(b) $n_1 = 250$ first-stage variables

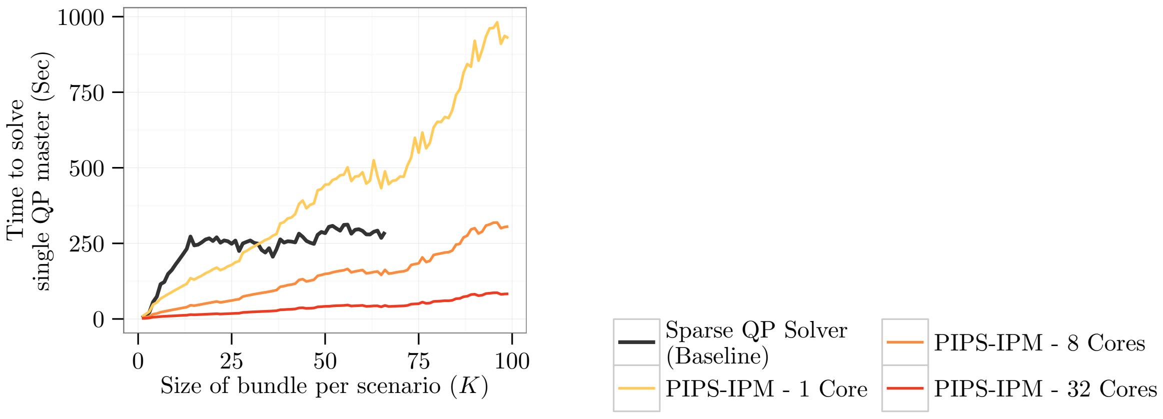

(c) $n_1 = 1,500$ first-stage variables

Figure 3: Time to solve QP master for small, medium, and large `prod` instances, with $r = 1,000$ scenarios each. Black line terminates when the sparse QP solver (OOQP) failed due to out-of-memory error. PIPS-IPM is a parallel block-angular-structure-exploiting QP solver applied to (22). Number of first-stage variables and size of the bundle (number of columns) per scenario primarily determine the difficulty of the QP to solve. Parallel speedups at the largest bundle size for (1 to 8 cores, 1 to 32 cores) are (8.5x, 22x), (8.8x, 28x), and (3.0x, 11x), for the small, medium, and large instances, respectively.

and computational, is required to address the load imbalance, perhaps considering asynchronicity akin to the work of Linderoth and Wright [18]. This is in addition, of course, to implementing a branching scheme.

In light of the recent availability of affordable multicore architectures and on-demand distributed computing, parallelizable optimization algorithms such as dual decomposition have the potential to be widely used if they can be shown to provide significant speedups over the state of the art on a single desktop machine.

Our analysis has been limited to two-stage formulations; we leave a treatment of multistage stochastic programming for future work. We remark, however, that the *nested* block-angular structure that could arise in the master of a multistage problem remains within the framework of parallel structure-exploiting interior-point methods; in particular, nested structure was considered by [26].

Acknowledgments

This work was supported in part by the U.S. Department of Energy under Contract DE-AC02-06CH11357. This research used resources of the Laboratory Computing Resource Center at Argonne National Laboratory, which is supported by the Office of Science of the U.S. Department of Energy under contract DE-AC02-06CH11357. Financial support from the University of Chicago Booth School of Business is also gratefully acknowledged. We thank Christoph Helmberg for correspondence regarding the `ConicBundle` package.

[1] J. R. Birge, F. Louveaux, Introduction to Stochastic Programming, 2nd Edition, Springer, New York, 2011.

[2] S. Ahmed, Two-stage stochastic integer programming: A brief introduction, in: J. J. Cochran, L. A. Cox, P. Keskinocak, J. P. Kharoufeh, J. C. Smith (Eds.), Wiley Encyclopedia of Operations Research and Management Science, Wiley, 2010.

[3] S. Sen, Algorithms for stochastic mixed-integer programming models, in: K. Aardal, G. L. Nemhauser, R. Weismantel (Eds.), Discrete Optimization, Elsevier, 2005, pp. 515–558.

[4] C. C. Carøe, R. Schultz, Dual decomposition in stochastic integer programming, Operations Research Letters 24 (1-2) (1999) 37–45.

[5] G. Lulli, S. Sen, A branch-and-price algorithm for multistage stochastic integer programming with application to stochastic batch-sizing problems, Management Science 50 (6) (2004) 786–796.

[6] A. Frangioni, About Lagrangian methods in integer optimization, Annals of Operations Research 139 (1) (2005) 163–193.

[7] C. Lemaréchal, Lagrangian relaxation, in: M. Jünger, D. Naddef (Eds.), Computational Combinatorial Optimization, Springer, 2001, pp. 112–156.

[8] K. Anstreicher, L. Wolsey, Two "well-known" properties of subgradient optimization, Mathematical Programming 120 (1) (2009) 213–220.

[9] J. B. Hiriart-Urruty, C. Lemaréchal, Convex Analysis and Minimization Algorithms, Vol. I-II, Springer-Verlag, Germany, 1993.

[10] K. Jörnsten, M. Näsberg, P. Smeds, Variable splitting: A new Lagrangean relaxation approach to some mathematical programming models, Tech. Rep. LiTH-MAT-R-85-04, University of Linköping (1985).

[11] K. C. Kiwiel, Approximations in proximal bundle methods and decomposition of convex programs, Journal of Optimization Theory and Applications 84 (3) (1995) 529–548.

[12] S. Feltenmark, K. C. Kiwiel, Dual applications of proximal bundle methods, including

Lagrangian relaxation of nonconvex problems, SIAM Journal on Optimization 10 (3) (1999) 697–721.

[13] W. Römisch, S. Vigerske, Recent progress in two-stage mixed-integer stochastic programming with applications to power production planning, in: P. M. Pardalos, S. Rebennack, M. V. F. Pereira, N. A. Iliadis, P. M. Pardalos (Eds.), Handbook of Power Systems, Vol. I, Springer, 2010, pp. 177–208.

[14] A. Ruszczyński, A regularized decomposition method for minimizing a sum of polyhedral functions, Mathematical Programming 35 (3) (1986) 309–333.

[15] R. E. Marsten, W. W. Hogan, J. W. Blankenship, The boxstep method for large-scale optimization, Operations Research 23 (3) (1975) 389–405.

[16] C. Lemaréchal, A. Nemirovskii, Y. Nesterov, New variants of bundle methods, Mathematical Programming 69 (1) (1995) 111–147.

[17] K. C. Kiwiel, Proximal level bundle methods for convex nondifferentiable optimization, saddle-point problems and variational inequalities, Mathematical Programming 69 (1-3) (1995) 89–109.

[18] J. Linderoth, S. Wright, Decomposition algorithms for stochastic programming on a computational grid, Computational Optimization and Applications 24 (2) (2003) 207–250.

[19] C. I. Fábián, Z. Szöke, Solving two-stage stochastic programming problems with level decomposition, Computational Management Science 4 (4) (2007) 313–353.

[20] W. Oliveira, C. Sagastizábal, S. Scheimberg, Inexact bundle methods for two-stage stochastic programming, SIAM Journal on Optimization 21 (2) (2011) 517–544.

[21] O. Briant, C. Lemaréchal, P. Meurdesoif, S. Michel, N. Perrot, F. Vanderbeck, Comparison of bundle and classical column generation, Mathematical Programming 113 (2) (2008) 299–344.

[22] V. Zverovich, C. Fábián, E. Ellison, G. Mitra, A computational study of a solver system for processing two-stage stochastic LPs with enhanced benders decomposition, Mathematical Programming Computation 4 (3) (2012) 211–238.

[23] K. C. Kiwiel, Proximity control in bundle methods for convex nondifferentiable minimization, Mathematical Programming 46 (1) (1990) 105–122.

[24] G. M. Amdahl, Validity of the single processor approach to achieving large scale computing capabilities, in: Proceedings of AFIPS spring joint computer conference, ACM, Atlantic City, NJ, 1967, pp. 483–485.

[25] J. Gondzio, R. Sarkissian, Parallel interior point solver for structured linear programs, Mathematical Programming 96 (3) (2003) 561–584.

[26] J. Gondzio, A. Grothey, Parallel interior-point solver for structured quadratic programs: Application to financial planning problems, Annals of Operations Research 152 (1) (2007) 319–339.

[27] K. C. Kiwiel, A Cholesky dual method for proximal piecewise linear programming, Numerische Mathematik 68 (3) (1994) 325–340.

[28] J. Forrest, CLP.
URL `https://projects.coin-or.org/Clp`

[29] C. Helmberg, ConicBundle.
URL `http://www-user.tu-chemnitz.de/~helmberg/ConicBundle/`

[30] E. M. Gertz, S. J. Wright, Object-oriented software for quadratic programming, ACM Transactions on Mathematical Software 29 (1) (2003) 58–81.

[31] I. S. Duff, MA57: A code for the solution of sparse symmetric definite and indefinite systems, ACM Transactions on Mathematical Software 30 (2) (2004) 118–144.

[32] M. Lubin, C. G. Petra, M. Anitescu, V. Zavala, Scalable stochastic optimization of complex energy systems, in: Proceedings of 2011 International Conference for High Performance Computing, Networking, Storage and Analysis, ACM, Seattle, WA, 2011, pp. 1–64.

[33] E. Anderson, Z. Bai, C. Bischof, S. Blackford, J. Demmel, J. Dongarra, J. Du Croz, A. Greenbaum, S. Hammarling, A. McKenney, D. Sorensen, LAPACK Users' Guide, 3rd Edition, SIAM, Philadelphia, PA, 1999.

[34] S. Mehrotra, On the implementation of a primal-dual interior point method, SIAM Journal on Optimization 2 (4) (1992) 575–601.

[35] T. Achterberg, SCIP: Solving constraint integer programs, Mathematical Programming Computation 1 (1) (2009) 1–41.

[36] W. Gropp, R. Thakur, E. Lusk, Using MPI-2: Advanced Features of the Message Passing Interface, 2nd Edition, MIT Press, Cambridge, MA, 1999.